Interferometric computation beyond quantum theory

Andrew J. P. Garner*

Centre for Quantum Technologies, National University of Singapore,
3 Science Drive 2, 117543, Singapore

(Dated: October 17, 2016)

There are some computational problems that admit a quantum solution more efficient than the fastest known classical alternatives. One explanation for this efficiency is that quantum theories allow us to enter every possible input of a function as a coherent superposition, and perform a calculation by way of interference: an *interferometric computation*. There has been recent foundational interest in theories of reality that go beyond quantum theory. Here, we present a generalized formulation of interferometric computation, and explore how theories can differ from quantum theory and still perform distributed calculations in this set-up. We shall see that *quaternionic quantum theory* proves a suitable candidate, whereas *box-world* does not. We also find that a classical hidden variable model first presented by Spekkens[a] can only be used for this type of computation when the system is subject to an epistemic restriction.

arXiv:1610.04349v1 [quant-ph] 14 Oct 2016

I. INTRODUCTION

There are algorithms available for quantum computers that solve particular tasks faster than any known deterministic classical algorithm [1–7]. It is argued that the quantum efficiency arises from the ability to prepare a superposition of input bit-strings, and *interfere* them using the quantum computer in order to produce a meaningful output [6]: a process that is not possible with a classical computer.

There has been much foundational interest in theories of reality that go beyond quantum theory, perhaps incorporating quantum theory as a limiting case, or even replacing it in yet-unencountered regimes. Often, one considers a set of fundamental principles (*axioms*) and uses these to derive quantum theory (see e.g. [8–13]). By testing these foundational principles one can validate quantum theory as a good description of reality; by relaxing them one can create *foil theories* (alternatives) that exhibit contrasting behavior to quantum theory, which might be observable in experimental tests.

In particular, one class of post-quantum theories known as *box-world* concerns itself with theories that are more non-local than quantum theory [14–18]. If such non-local states existed, it has been shown that they would make certain tasks involving distributed computing trivial [19, 20]. This unusual computational property of a non-quantum theory motivates us to ask if interferometric computation (which is possible in quantum theory) can be extended into other theories. By building on recent work that extends the concept of interference beyond quantum theory [21–26], we shall address this question in this article .

In particular, we will focus on the interferometric implementation of the Deutsch–Jozsa algorithm [1, 2, 6] and Grover's algorithm [4, 7]. We will discuss how to represent this type

[a] Phys. Rev. A, 75:3:32100, 2007

* ajpgarner@nus.edu.sg

of algorithm in a generalized setting. We will then show that they can not be implemented in box-world, but can be extended into quaternionic quantum theory. Finally, we will consider the case of Spekkens' toy model [27, 28], a classical hidden variable theory, which can be made to exhibit certain aspects of quantum behavior by imposing a restriction on the allowed states of knowledge of the system. In this case, we find that only with the epistemic restriction is the computation possible, even though non-trivial dynamics are possible in both cases.

II. TECHNICAL INTRODUCTION TO POST-QUANTUM THEORIES

The GPT framework.— We will find it convenient to adopt the framework of generalized probabilistic theories (GPTs) as presented for instance in [9, 11, 29]. This framework amounts to taking the operational view that a system can be described in terms of the measurements made on it, and the outcomes associated with any given measurement. The states of the system are then a collection of probability distributions over these measurements. It is assumed that by knowing the statistics of some finite number of different measurements (an *informationally complete set*), one can predict the outcomes of *any* measurement made on the system. These statistics are known as the *state* of the system, which can thus be written as a real-valued vector . Likewise, there is a natural dual space of measurement outcome *effects*, which are functionals that act on states to give the probability associated with any particular measurement outcome. Effects are written as real-valued (co-)vector e, such that for some outcome i associated with effect e_i, the probability that this outcome occurs for a state s is given by the inner product $P(e_i|s) = e_i \cdot s$.

For example, consider a (classical) coin which has either "heads" or "tails" facing upwards. The state of the coin may be written as $s = [P(\text{heads}), P(\text{tails})]^{\mathrm{T}}$, and the effects associated with heads and tails are $e_{\mathrm{H}} = (1, 0)$ and $e_{\mathrm{T}} = (0, 1)$ respectively. By inspection, $e_{\mathrm{H}} \cdot s = P(\text{tails})$. Following a fair toss, the state would be written $s = \left(\frac{1}{2}, \frac{1}{2}\right)$ and the probability of being in heads (or tails) is obviously given by $\frac{1}{2}$.

A more involved example would be to describe the state of the elementary two-level quantum system: the qubit. By knowing the outcome probabilities of measurements associated with the Pauli matrices (σ_x, σ_y and σ_z), one can fully recreate the state (i.e. density matrix in the quantum formalism). Thus, within the GPT framework, one can express the qubit state as a real positive vector:

$$s = \begin{pmatrix} P(X=+1) \\ P(X=-1) \\ \hline P(Y=+1) \\ P(Y=-1) \\ \hline P(Z=+1) \\ P(Z=-1) \end{pmatrix}. \tag{1}$$

In quantum theory, there are restrictions on the probability distributions for the measurements within a state. For instance, a state where $P(Z=+1) = 1$ and $P(X=+1) = 1$ is disallowed because it violates the uncertainty principle. The uncertainty restriction is very much a property of quantum theory (following from the Cauchy-Schwarz inequality applied to a Hilbert space [30]). Moreover, because the set of Pauli matrices form a mutually-unbiased basis, if any measurement has an outcome which occurs with certainty

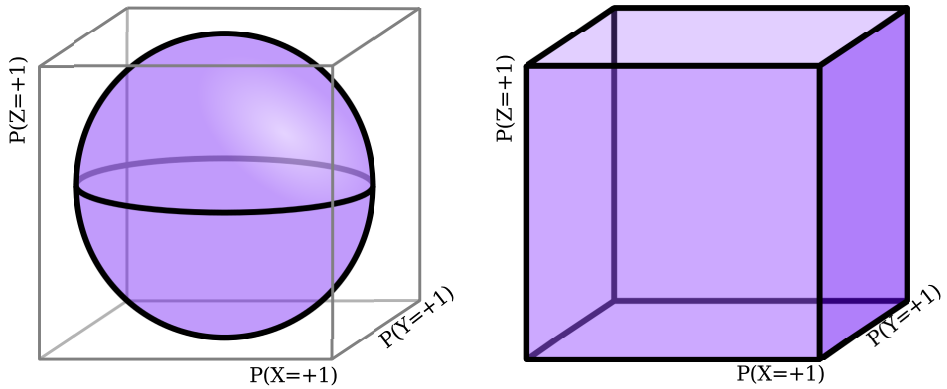

FIG. 1: ***Qubit and gbit state-spaces.*** *LHS shows the allowed states a qubit can take. RHS shows the allowed states a (three-measurement) gbit can take. The uncertainty principle prevents the qubit from taking arbitrary states; as a consequence,*

(probability 1), then the outcomes for all other measurements occur according to a uniformly random distribution. Geometrically, the uncertainty principle on a qubit describes a spherical state-space (set of valid states), as illustrated on the LHS of fig. 1.

It is possible to conceive of alternative theories where this restriction does not hold. *Gbits*, which form the single-system state-spaces in *box-world*, admit any valid probability distribution (subject to the requirement that all measurements normalize to the same value). A gbit composed of three binary measurements could be written in exactly the same form as the quantum bit (eq. (1)), but without the uncertainty principle, its state-space is given by a cube, as illustrated on the RHS of fig. 1.

Finally, we remark that transformations in the GPT framework are linear maps between allowed states. For classical theories (where there exists a single informationally complete measurement), these transformations are exactly the set of stochastic matrices acting on a probability vector, whereas in general they have a slightly more complicated form (see *Theorem 7* in [29]). In all cases, they are real-valued matrices. Reversible transformations take the state-space to itself and are thus the automorphism group of the state-space (or some subset thereof if there are other restrictions imposed on the theory). For example, single-qubit reversible transformations are given by $\mathrm{SO}(3) \subset \mathrm{O}(3)$, whereas gbit reversible transformations correspond to the finite group formed by relabelling measurements and measurement outcomes (e.g. for three binary measurements, this gives the octahedral group of cubic symmetries).

Spekkens' toy theory.— Spekkens has presented a classical hidden-variable toy theory that replicates many of the features of quantum theory [27, 28], through the introduction of an *epistemic restriction* that restricts states to only contain half the amount of knowledge required to fully determine the hidden variable. The elementary system in this model consists of four classical possibilities (i.e. two bits), of which only one bit of information can be known. We can pictorially represent the four ontic possibilities as , , and .

The valid single-bit epistemic states thus consist of two ontic possibilities, and can be divided into three groups of mutually-exclusive pairs: { , }, { , } and {, } which by analogy with quantum theory we label as the ± 1 outcomes of X, Y and Z respec-

tively. These pairs of states are naturally associated with measurements that ask which of the two options is consistent with the underlying ontic state. Following any measurement, the ontic state is moved to one of the supported possibilities with equal probability (preventing the violation of the epistemic principle by making two different measurements in a row). From this, it follows that if a measurement is made on an epistemic state which itself is not one of the outcomes of the measurement (e.g. measurement vs. made on the state), then either outcome should occur with equal probability.

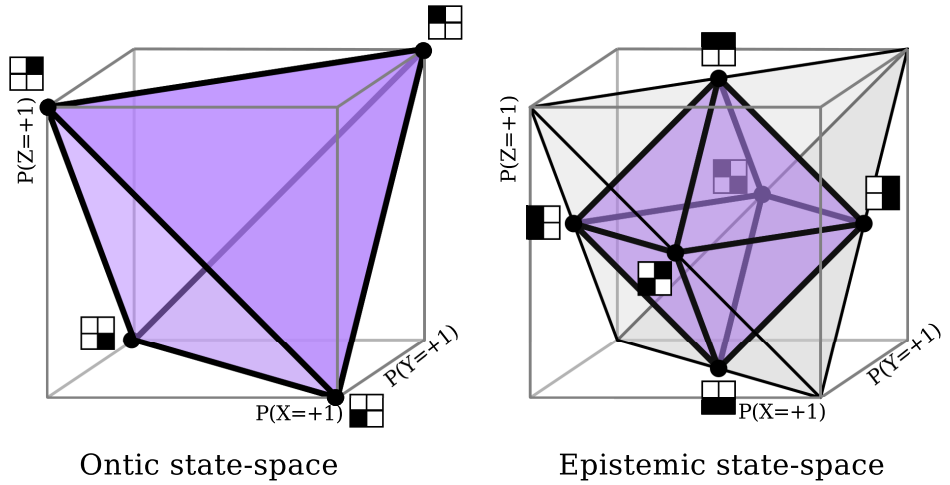

FIG. 2: ***State-spaces of Spekkens' toy theory.*** *LHS shows the* ontic *tetrahedron of states that the hidden variable could take. RHS shows the* epistemic *octahedron of states that satisfy the epistemic restriction.*

Although this theory not *a priori* part of the GPT framework, it can be naturally extended by taking the convex combination of the statistics of these states [24] (cf. other treatments in [31, 32]). The states may then be written in terms of the X, Y and Z statistics, and represented by a vector of the form eq. (1), making it naturally comparable with qubits and three-measurement gbits (sharing the same effects for these measurements). The set of allowed states can be plotted on Cartesian axes, as shown in fig. 2. The statistics spanned without an epistemic restriction form a tetrahedron, whereas with the epistemic restriction, one has an octahedron (identical to the convex hull of one-bit stabilizer states [33]). Because transformations in Spekkens' toy theory correspond to permutations on the hidden variable, for a single bit they are given by the automorphism group of the tetrahedron S_4; not all of the symmetries of the octahedron are allowed.

III. INTRODUCTION TO GENERALIZED PHASE AND INTERFERENCE

In this paper, we shall consider phase as presented in [24–26] (cf. the extension of Sorkin's hierarchical family of interference patterns [21–23]). Phase between possibilities is one of the most remarkable features of quantum theory. Much as how knowing the amplitude of a classical wave is not sufficient to predict how two waves interfere, simply knowing the probability associated with a particular possibility is not in general sufficient to predict the behavior of a quantum system. For example, consider the qubit states $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$

and $|-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$, when measured in the computational basis ($|0\rangle$ vs. $|1\rangle$), both states are equally likely to give either outcome. However when a Hadamard gate $H = \frac{1}{\sqrt{2}}\left(\begin{smallmatrix} 1 & 1 \\ 1 & -1 \end{smallmatrix}\right)$ is applied to the two states: $H|+\rangle = |0\rangle$ and $H|-\rangle = |1\rangle$, and so afterwards the two states would give opposite answers in the computation basis with certainty.

For quantum theory in general, for some basis $\{|\xi_j\rangle\}$, all states of the form $|\psi\rangle = \sum_j a_j e^{i\phi_j}|\xi_j\rangle$ for $a_j, \phi_j \in \Re$ with fixed values of $\{a_j\}$ will have the same statistics with respect to the $\{|\xi_j\rangle\}$ measurement for all values of $\{\phi_j\}$. Moreover, any transformation of the form $U_\xi = \sum_j e^{i\alpha_j}|\xi_j\rangle\langle\xi_j|$ will not alter the $\{|\xi_j\rangle\}$ statistics. Transformations of this kind are known as *phase operations*, and in quantum theory, it is as pertinent to consider the set of of allowed unitary matrices of this form as it is the set of allowed values of $\{\phi_i\}$ in $|\psi\rangle$, as one can generate a state with arbitrary $\{\phi_i\}$ values from some reference state, and free choice of the operation U_ξ.

This concept can be extended to the GPT framework as follows [24]:

Definition 1 (Phase operation)**.** A transformation T is a *phase operation* associated with a measurement M if it never alters the statistics associated with M, such that:

$$e_i \cdot s = e_i \cdot Ts \tag{2}$$

for every effect e_i associated with M, and every state s admitted in the theory.

In the case of the reversible transformations, the phase restriction induces a sub-group of transformations which we refer to as the *phase group*.

As we shall be interested in *branching interferometers* (e.g. the Mach–Zehnder interferometer in fig. 3), we need an operational method of identifying whether a particular transformation can be implemented on a particular branch. If we imagine some element that might be added to a branch (for instance, a piece of glass altering the optical path-length) which does not cause the particle to jump between branches, then this naturally corresponds to an element of the phase group of the system with respect to the "which branch?" measurement (we will label this Z by convention). However, it is not *a priori* obvious if every transformation in this group can be induced by an action taken local to one particular branch. In quantum theory, one might by convention associate

$$U_i = e^{i\Phi}\left(e^{i\phi_i}|i\rangle\langle i| + \sum_{j\neq i}^{N}|j\rangle\langle j|\right) \tag{3}$$

with an action taken on branch i, which when spanned across all ϕ_i in one branch forms a subgroup of the $\bigoplus^N U(1)$ phase group. This form guarantees that an action taken on, say, branch 1 does not induce a phase difference between remote branches 2 and 3; only for a two-branch system are the two groups are identical, due to the global U(1) phase freedom. As the above form does not obviously generalize beyond Hilbert spaces, we will use the following operational concept to identify the local subgroups [25]:

Definition 2 (Branch locality)**.** A transformation T_i can be localized to branch i if it does not alter any states that have no probability of being in branch i. For position measurement effect z_i associated with branch i, where $s_{\neg i}$ is an arbitrary state with no support on branch i:

$$e \cdot s_{\neg i} = e \cdot T_i s_{\neg i} \quad \forall e,\ s_{\neg i} \in \{s \,|\, z_i \cdot s = 0\} \qquad \leftrightarrow \qquad T_i s_{\neg i} = s_{\neg i}. \tag{4}$$

Any T that satisfies this for branch i may be said to be *localizable* to branch i; though this does not imply that T can *only* be localized to i; for example, the identity element $\mathbb{1}$ satisfies branch locality for all branches in every theory. It is shown in appendix A that definition 2 implies the conventional quantum form written in eq. (3).

This is a *locality* restriction, because it captures an intrinsic quality of something being "over here" instead of "over there" (cf. "non-locality" in the context of CHSH-Bell violation [14–17], which concerns itself with the impossibility of quantum *local realism*). Without branch locality, one would find it conceptually very difficult to ever isolate an experiment from the rest of the universe: every action taken *everywhere* would have to be taken into account in order to predict the outcomes of an experiment. Thus, we will axiomatically impose branch locality as a reasonable restriction on the behavior of any spatially-disjoint interferometer.

IV. INTERFEROMETRIC COMPUTATION: CONSTANT VS. BALANCED

Deutsch–Jozsa algorithm.— Suppose there is some function $f : \{0,1\}^{\otimes n} \to \{0,1\}$, which is guaranteed either to be *constant* such that $f(x)$ for some bit-string $x \in \mathcal{B}^n = \left\{\{0,1\}^{\otimes n}\right\}$ evaluates to the same value for all choices of input, or *balanced* such that $f(x) = 1$ for exactly half of the inputs (and 0 for the other half). Classically one might need to test f with up to $2^{n-1}+1$ inputs before it could be verified to be constant (a balanced function could be discovered after 2 inputs, but might look constant for the first 2^{n-1} tests). However, by inputting a quantum bit-string into the device, the Deutsch–Jozsa algorithm [1, 2, 6] can verify whether the function is constant or balanced in just one query. The efficient quantum encoding [6] involves $n+1$ qubits, and a "quantum oracle" which performs an f-controlled-NOT gate that evaluates f on the first n qubits and applies a NOT gate on the final qubit if $f(x) = 1$. By preparing all of the qubits in the $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ state, the final qubit remains as $|+\rangle$ if the function is constant and goes to $|-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$ if it is balanced.

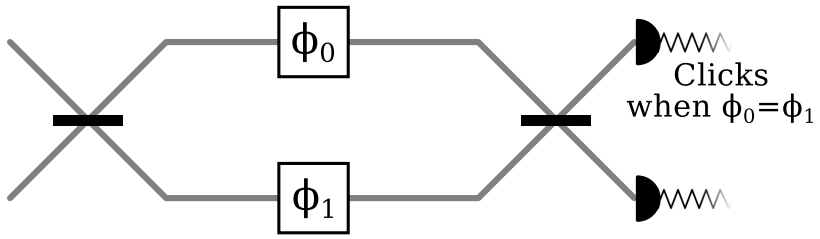

FIG. 3: ***Mach–Zehnder interferometer.*** *A photon incident from the upper branch is put into a superposition of both branches by the first beamsplitter. Phases ϕ_0 and ϕ_1 are induced on each arm. The branches recombine at the second beamsplitter, and if $\phi_0 = \phi_1$, then the photon finishes in the upper branch with certainty. If $|\phi_0 - \phi_1| = \pi$, then the photon finishes in the lower branch with certainty.*

There is also a simple single-photon interferometry experiment that solves the fixed vs. balanced problem using one large degree of freedom. Suppose each path encodes a different bit-string ($x \in \mathbb{B}^n$). A beamsplitter (or series of beamsplitters) can be used to prepare a coherent superposition of all possible bitstrings, written as $\left(\frac{1}{\sqrt{2}}\right)^n \sum_{x\in\mathbb{B}^n} |x\rangle$. The function

f can then be encoded as a collection of phases induced on branch x such that if $f(x) = 0$ then a phase 0 is applied to branch $|x\rangle$, and if $f(x) = 1$ then a phase of π is applied. Thus, the unitary phase operation $\sum_{x\in\mathbb{B}^n} e^{i\pi f(x)}|x\rangle\langle x|$ is applied over all 2^n branches. Finally, a set of beamsplitters is used to execute the operation $H^{\otimes n}$. If f is constant, then the photon will always be found in branch 0, and if it is balanced the the photon will never be found in branch 0. By placing a detector in this branch, it is possible to tell if f is balanced or constant in a single run of the experiment. The single-bit version of this protocol is exactly implemented by a Mach–Zehnder interferometer (fig. 3).

Such an implementation is arguably not as efficient as the Deutsch–Jozsa algorithm, as it requires 2^n branches (and hence 2^n optical components), but does not require any entangling gates. However, when treated as an oracle problem, this solution only requires one oracle query: i.e. can be performed by passing just one single photon through the system, and so if transmission of data to and from the oracle is the expensive part of the task (say because half of the experiment is very far away, or the oracle charges per access [7]), then this could still be viewed as an efficient solution with respect to oracle calls.

In this article, we shall focus on the interferometric algorithm, as it can be considered purely in the context of phase and interferometry, in a way that naturally generalizes beyond quantum theory, and allows for interesting exploration of post-quantum dynamics.

Interferometric computation in generalized probabilistic theories.— Let us now address the interferometric computation of the constant vs. balanced problem in generalized probabilistic theories. We will make the assumption that all theories we consider are *operational* such that there is a method of preparing any state in the state-space, and that all statistics required to specify the state of the system can be extracted by some kind of measurement (if the experiment is repeated enough times). This broadly amounts to abstracting away the initial and final beamsplitters (as well as the measurement device) and instead focusing our attention on the central region where the branches are spatially disjoint.

In order to execute the algorithm (for n bits over 2^n branches), we suggest that the following criteria are required:

i. For every $x \in B^n$, there must be at least two choices of transformations T^1_x and T^0_x that can be applied locally to branch x, which may be associated with $f(x) = 1$ and $f(x) = 0$ respectively.

ii. There must be some effect e_{C} such that when the output state s_{out} is measured, $e_{\mathrm{C}}\cdot s_{\mathrm{out}} = 1$ if f is constant, and $e_{\mathrm{C}} \cdot s_{\mathrm{out}} = 0$ if f is balanced, for at least one input state s_{in} chosen independently of f.

Requirement i frames the set-up in the context of a *distributed computation*: one could imagine 2^n independent agents, each of which who calculates $f(x)$ locally, and chooses a phase plate to add to their branch based on the result of their calculation. The transformation induced by the phase-plate put in by agent x does not depend on the outcomes of $f(y \neq x)$ obtained by the other agents, but might be dependent on which branch the agent is in, allowing all of the values of $f(x)$ to be computed in parallel. For example, in quantum theory if $f(00) = 1$, Agent 00 might put in a piece of glass that induces the operation $U^1_{00} = -|00\rangle\langle 00| + \sum_{i\in\{01,10,11\}} |i\rangle\langle i|$, whereas when $f(11) = 1$, Agent 11 puts in a piece of glass inducing the different operation $U^1_{11} = -|11\rangle\langle 11| + \sum_{i\in\{00,01,10\}} |i\rangle\langle i|$. Crucially, we require that the transformation induced by each agent is compatible with *branch locality* on that branch.

Requirement ii captures the deterministic-interferometry aspect of the algorithm: there must be a single measurement that determines whether f is constant or balanced. A slightly looser version would require there to be a e_{C} such that $e_{\mathrm{C}} \cdot s_{\mathrm{out}} > \frac{1}{2}$ if f is constant, and $< \frac{1}{2}$ if f is balanced, such that the two classes of f may ultimately be distinguished by enough repetitions of the experiment.

Quantum theory.— As expected, quantum theory is compatible with both requirements. In general, an agent x can implement $U_x^1 = -|x\rangle\langle x| + \sum_{i \neq x} |i\rangle\langle i|$ if $f(x) = 1$ and $U_x^0 = \sum_{i \in \mathbb{B}^n} |i\rangle\langle i| = \mathbb{1}$ otherwise. Transformations of this form are compatible with branch locality (proven explicitly in appendix A), and can be composited in any order (the operations are all diagonal in the same basis) such that the overall effect of all the agents is to induce the operation $U_f = \sum_{x \in \mathbb{B}^n} e^{i\pi f(x)} |x\rangle\langle x|$. By preparing the initial state $|+\rangle^{\otimes n}$, one can perform a final measurement in the $X \otimes \cdots \otimes X$ basis (i.e. by applying $H^{\otimes n}$ followed projecting into $|0 \ldots 0\rangle$). This has probability of a positive outcome given by $p = \frac{1}{2^n} |\sum_{x \in \mathbb{B}^n} e^{i\pi f(x)}|^2$. As $f(x) = 0$ contributes a term $e^0 = 1$ and $f(x) = 1$ contributes $e^{i\pi} = -1$, we see that $p = 1$ when f is constant as the global sign does not matter, and $p = 0$ for any balanced f because half the addends will cancel out the other half.

Classical theory.— On the other hand, classical theory does not permit this algorithm, but rather fails at requirement i since it does not admit any non-trivial (i.e. other than $\mathbb{1}$) phase transformations [24], and so $T_{\mathrm{x}} = \mathbb{1}$ always. There will therefore be no way to encode any measurable information about f through actions taken on the separate branches, and so the final measurement will only depend on the initial state s_{in}.

Box-world.— By these criteria, we can also confirm the suggestion made in [25] that such an interference algorithm can not be implemented in box-world as a distributed computation[1] (failing criterion i) when the position measurement is taken to be one of the gbit measurements that allows for complete freedom in selecting the other measurement statistics. This is because *no* transformation can be localized to *any* subset of branches (proven for the general case in [25]) without violating branch locality in some way, and so the agents on the separate branches have no way of encoding the outcome of $f(x)$.

Let us see this explicitly for the $n = 1$ case, where the "which-branch?" measurement is a binary degree of freedom and the family of gbits are d-dimensional hypercubes. As $d = 1$ is just classical theory, we shall consider the simplest non-trivial example, $d = 2$. The state-space of this theory is given by a square (sometimes called a *square bit*), and the phase-group is Z_2, consisting of either doing nothing ($\mathbb{1}$), or exchanging the statistics of the complementary measurement (X-flip). Writing the position and complementary measurements as Z and X, we note that any operation branch local to $Z = 1$ must not change states where $P(Z = 1) = 0$. However, as the action of X-flip is independent of Z, there is a state where $Z = -1$ and $X = 1$ which will be turned into one where $Z = -1$, $X = -1$ when the X-flip is applied. Thus, the X-flip operation can not be applied branch locally at $Z = 1$. Similar logic rules it out localization to $Z = -1$, and hence we see that although X-flip is in the global phase group, it can not be assigned to either branch. The only remaining group element is $\mathbb{1}$, and this does nothing, and hence the localizable phase dynamics are trivial.

If it were possible to apply $\mathbb{1}$ and X-flip from each branch, then the algorithm could be performed as an interferometric computation. However, as they can not, such a calculation can only be done by global phase, in which $\mathbb{1}$ is applied globally if the function is constant,

[1] It should also be noted that implementing the efficient Deutsch–Jozsa algorithm to solve this problem on a collection of gbits can also not be done in box-world, because the f-controlled-NOT gate will not be an allowed transformation, according to the restrictions on correlating dynamics shown in [35]

and X-flip if it is balanced (and then $e_C = x_+$ such that $x_+ \cdot s = P(X = +1)$ could satisfy criterion ii). This effectively reduces the problem to a trivial oracle that answers the question "is the function constant or balanced?" and no longer has the structure that allows the problem to be treated as a distributed computation.

That there is a global implementation, but not a local one, highlights a key difference between interference in box-world and interference in quantum theory. Processes that induce phase shifts in an intrinsically non-branch-local way (such as by the Aharonov–Bohm effect [34] in which switching on a magnetic solenoid between two branches induces a geometric phase) may still be plausible in box-world. In quantum theory, any phase operation induced by the Aharonov–Bohm effect could be equivalently induced by pieces of glass placed in each branch: there are no operations in the phase group $\mathcal{G}$ (in this case $\bigoplus^N \mathrm{U}(1)$) that are not in the union of all the operations local to each branch $\bigcup_i \mathcal{G}_i$. On the other hand, in box-world $\bigcup_i \mathcal{G}_i = \{\mathbb{1}\}$, and so even if there is a process that can induce a non-trivial global phase (such that $\mathcal{G} \neq \{\mathbb{1}\}$) there is no way to generate this phase from local operations. Thus, whilst quantum theory could do the algorithm distributing f over many branches, the best that is admissible in box-world would constitute of a single agent who evaluates all values of f to see if it is constant or balanced, and then switches on a solenoid only if it is constant.

d-balls and quaternionic quantum theory.— A d-ball is the generalization of the single-qubit Bloch-sphere into d dimensions, by generalizing the three binary measurements into d such measurements (see e.g. [13, 36–38]). For states with d measurements $\{X_i\}$, the state-space satisfies a quantum-like uncertainty relation given by $\sum_i \left(P(X_i = 1) - \frac{1}{2}\right)^2 \leq \frac{1}{4}$, such that if any measurement has a certain outcome, all others are uniformly random. Geometrically, this corresponds to a d-dimensional hypersphere. Excluding reflections (by analogy with quantum theory), the automorphism group of a single d-ball state-space is given by $\mathrm{SO}(d)$, and the phase group after fixing the values of one measurement is $\mathrm{SO}(d-1)$.

Let us first consider single-bit systems. First, we note that *all* operations in the phase group can be implemented locally on either branch. Intuitively, this is because for each branch, there is a unique state that has no support on that branch, and as this state has uniformly random outcomes for every measurement except the position one, it is invariant under application of $\mathrm{SO}(d-1)$, and thus will never violate the requirement of branch locality. Ignoring $d = 1$ as it is classical, and $d = 2$ (because it only has a non-trivial phase group if reflections are permitted), we see that for $d > 2$, we can pick a π (generalized) rotation about the Z axis and each agent can apply the same rotation when $f(x) = 1$ (and do nothing when $f(x) = 0$). The algorithm will then work exactly as the 1-bit quantum case did; and there will be a measurement whose outcome changes from one value to another if $f(0) \neq f(1)$. (As we only need the identity and one other operation, which forms an Abelian subgroup isomorphic to integer addition modulo 2 (Z_2), the issues of commutativity addressed in [38] should not cause problems in the implementation of this algorithm.)

The general algorithm is trickier to discuss, as the spherical shape of the state-space is specific to a two-level system (and thus generalizations of this will not provide obvious insight as to what measurements and transformations are possible on a $N = 2^n$ level system). However, the d-balls are in general the state-spaces associated with two-level systems described by a Jordan algebra [39], which includes real-valued ($d = 2$), complex ($d = 3$) and quaternionic [41] ($d = 5$) quantum theory (i.e. variants of quantum theory where the Hilbert space is on a different field). Thus, at least in these cases, we might attempt to address the n-bit problem using the appropriate algebra. (By using the many-level single-system

approach, we also dodge the bullet of attempting to compose a multipartite system of n d-balls, where in general it is tricky to find the joint state-space since the tensor-product of such state-spaces is not uniquely defined.)

Let us then single out the example of quaternionic[2] quantum theory (whose two-level system is the 5-ball), which admits as its reversible operations elements of the symplectic group $\mathrm{Sp}(N)$ [41, 42]. As shown in shown appendix B, an N-level quaternionic system contains $\bigoplus^N \mathrm{Sp}(1)$ in its phase group (where $\mathrm{Sp}(1)$ are the unit quaternions; isomorphic to $\mathrm{SU}(2)$). Moreover each $\mathrm{Sp}(1)$ addend corresponds to operations that can be applied locally on a particular branch, and quaternionic states that different only by a sign are operationally indistinguishable (i.e. there is a Z_2 *global phase*, much like the $\mathrm{U}(1)$ complex quantum global phase).

In the context of evaluating f, each agent could thus choose between implementing $\mathbb{1}$ or $-1 \in \mathrm{Sp}(1)$ (where $-1 \cdot -1 = \mathbb{1}$ and corresponds to changing the quaternionic phase of a possibility by a real factor of -1) on their branch depending on the outcome of f. One can prepare the quaternionic superposition state $\frac{1}{\sqrt{N}} \sum_j |j\rangle$ (much like in complex quantum theory), and the action of the agents over the N branches will be to introduce phase of -1 in front of either all, none or exactly half of these terms. Because of the Z_2 global phase, the two balanced functions will result in operationally indistinguishable states; and in general, any function f could be replaced by its negation without changing the statistics output by the interferometer[3].

Finally, we must argue for the existence of a quaternionic beamsplitter whose effect is to sum together all these terms. This can be seen by noting that $\mathrm{Sp}(N)$ is isomorphic to $\mathrm{U}(N, \mathbb{H})$, which contains the usual quantum Hadamard gate that only contains real numbers and is sufficient for this task. (In general, there will be more types of quaternionic beamsplitter than quantum beamsplitter; but we only need one for this algorithm). Thus requirement ii is also satisfied, and we hence conclude that an interferometric computation to solve the constant/balanced problem can be done in quaternionic quantum theory.

Spekkens' toy model.— Let us consider the 1-bit implementation of the problem in the context of Spekkens' toy model. By comparison of the behavior with and without the epistemic restriction, we can see explicitly how the uncertainty principle makes the algorithm possible. The phase group reversible operations that preserve the Z measurement of Spekkens' toy model is $Z_2 \oplus Z_2 = \{1234, 2134, 1243, 2143\}$ (using the one-line permutation notation). By application of branch locality to the ontic state-space, we find that the upper branch can locally execute the operations $Z_2 \oplus \mathbb{1} = \{1234, 2134\}$ and the lower branch can locally execute $\mathbb{1} \oplus Z_2 = \{1234, 1243\}$ (and if both branches implement a non-trivial transformation the final element 2143 is performed). This may be seen visually in fig. 4.

The single-bit algorithm is effectively modulo-2 addition; but without the epistemic restriction, the operations on the two branches affect different parts of the group. Thus, there is no way for any action done by an agent on one branch to be canceled out by an action done by the agent on the other. There will be no way for constant and balanced f to result in mutually exclusive output states. For example, suppose was input, and each agent applies $\mathbb{1}$ if $f(x) = 0$, and the non-trivial element otherwise; the outcomes when the function

[2] Quarternions may be thought of as the extension of the field of complex numbers [40] to include three different imaginary elements i, j and k such that $ii = jj = kk = ijk = -1$. See appendix B and [41, 42].

[3] In general balanced functions which are not negations of each other result should result in operationally distinguishable statistics for some measurements, even if they all satisfy $e_C \cdot s = 0$. This can be conceptually understood by discarding half of the branches and performing interferometry on the remaining half: the choice of branches included in the subset in will general govern whether the subset is balanced, constant or neither, and if it was one of the first two, then a lower-dimension version of the algorithm should work on these. It would thus be disingenuous if all balanced functions resulted in the same statistics.

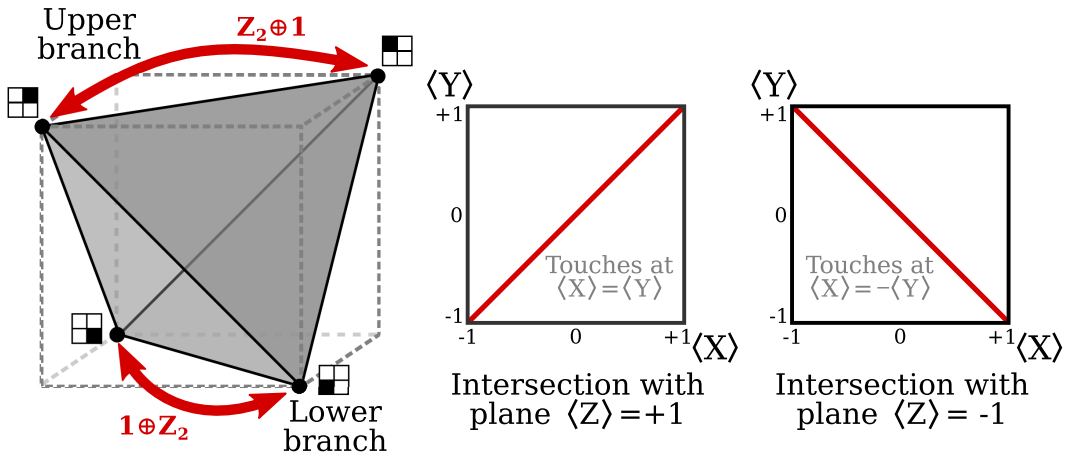

FIG. 4: ***Localisable dynamics for the ontic state-space of Spekkens' toy theory*** *The phase group $Z_2 \oplus Z_2$ splits into a different subgroup for each branch (draw on LHS). For the lower branch, there is a line of states with support on this branch: the states on the upper branch, drawn in the middle figure. Branch locality mandates that each of these states must not be altered by a transformation on the lower branch, and so one can either do nothing or exchange* ◨ *and* ◧, *resulting in the group $\mathbb{1} \oplus Z_2$. Similarly, for the upper branch, states with no support (RHS) must not be changed, and this allows only for doing nothing or exchanging* ◩ *and* ◪, *yielding the group $Z_2 \oplus \mathbb{1}$.*

is constant will be ◨ or ◧, and if the function is balanced will be ◩ or ◪. This failure is more subtle than that of box-world: here, we can find an encoding of $f(x)$ outcomes to operations on branch x that seemingly satisfies requirement i, but there is no such choice of encoding that allows an appropriate effect to be found that satisfies requirement ii. No single measurement can distinguish between these two cases, and one would have to perform complete tomography on the state (hence identifying f) before one could say if f were constant or balanced. This fails even the weaker version of requirement ii: no single effect has $e_C \cdot s > \frac{1}{2}$ if and only if f is constant. One might interpret this as there being a complete lack of global phase in the ontic state-space, such that the action on one branch can not cancel out with the action taken on another (cf. complex and quarternionic quantum theory where $\mathrm{diag}\,(-1, 1)$ and $\mathrm{diag}\,(1, -1)$ are equivalent to each other, since they only differ by a global phase).

This is no longer the case when we apply the epistemic restriction. The state-space is now subject to a quantum-like uncertainty principle such that if the outcome of Z is known with certainty, no degrees of freedom remain. This means that branch locality no longer divides the phase group into two separate subgroups, but rather admits for the application of any element in $Z_2 \oplus Z_2$ from either branch. Thus both agents could choose (for example) to execute 1234 when $f(x) = 0$ and 2143 when $f(x) = 1$. These operations will combine to form a representation of Z_2, and hence the single-bit Deutsch algorithm will be possible once more. For example, ◨ → ◧ when f is constant and ◨ → ◩ when f is balanced, which can be distinguished through a measurement of X.

In terms of the hidden variable, allowing for these operations leads to remarkable behaviour. If one were to interpret ◨ and ◩ as lying definitely with one agent and ◧ and ◪ with the other (reasonable under branch locality; see fig. 4) then with the epistemic re-

striction, an agent can seemingly violate branch locality on the hidden variable, and change the state of the system even when it has no support in her branch. However, because of the epistemic restriction, this violation can never be observed. A similar phenomenon can be observed in another famous hidden variable theory: Bohmian mechanics [43–45]. In Bohmian mechanics it is possible [45] to construct a set-up involving an entangled pair of particles entering Stern–Gerlach devices at two different positions in space, such that if the exact initial position of the one of the particles was known, at one detector it would be possible to determine from the final trajectory taken by the particle whether or not a measurement had made by the remote detector, violating no-signalling. However, because knowing the initial position is forbidden, this violation is avoided by an epistemic restriction.

V. UNORDERED DATABASE SEARCH: GROVER'S ALGORITHM

A classical search to find one record in an unordered database containing N entries requires on average $N/2$ queries, whereas Grover's algorithm [4] provides a quantum method which can determine this after $\mathcal{O}\left(\sqrt{N}\right)$ queries. Much like the Deutsch–Jozsa algorithm, one can implement the Grover algorithm without recourse to entanglement by building a sufficiently large single-system set-up [7]. In such a set-up, one passes a single particle through a series of beamsplitters such that it is in a superposition over all N branches. These N branches are connected to the database, which could be thought of as N black boxes, one containing a phase plate that induces a π phase shift, and the rest do not change the phase. The database could be thought of as a function $f(x) = \{0, 1\}$ where $f(x) = 1$ for only one value of x, and $f(x) = 0$ otherwise. Following this, the output branches are passed through a second beamsplitter, and an additional phase-shift of π is added to the first branch only. This entire process is repeated $\sqrt{N}$ times, after which it is highly likely ($p > \frac{1}{2}$) that the particle will be in the branch corresponding to where the π phase plate is.

Much like the many-branch implementation of fixed vs. balanced problem, a larger number of optical components are required to prepare the superposition states of an N level system. However, if the resource to be minimized is the number of times the database can be probed (e.g. because it is far away, or charges per access), then the N-branch implementation (as opposed to the qubit implementation) of this algorithm is just as efficient.

The unordered search addressed by Grover's algorithm is slightly less straightforward to classify than the constant vs. balanced problem. Certainly, we will still need requirement i to be true in order to build our database; in this case the function f is no longer constant or balanced, but rather guaranteed to only be 1 for one value of x. Thus, every theory which failed on this requirement (classical bits, boxworld) will also be unable to provide the coherent oracle required for Grover's algorithm. In the box-world case, as with the constant vs. balanced problem, one can only find admit a "global oracle" which trivially sets the state such that it flags the correct value of x satisfying $f(x) = 1$). On the other hand, quaternionic quantum theory admits all the components required (local phase-shifts and beamsplitters), which will work in the same way as in quantum theory, and so should allow an implementation of Grover's algorithm.

VI. SUMMARY

By treating spatial degrees of freedom in the generalized probabilistic theories, we have seen that interferometric computation can be extended beyond quantum theory. The usefulness of a theory was shown to be very much reliant on the structure of available dynamics. In particular, locality considerations dictate whether the process can be interpreted as a distributed computation performed by many agents. We ruled out box-world as the dynamics could not be localized to different branches, and so the only interferometric solution to either the constant vs. balanced or the unordered search problems would require a trivial oracle that directly answers the question.

On the other hand, in theories with interesting dynamics interferometric computation is once more possible. In complex and quaternionic quantum theory, which both contain a non-trivial unmeasurable global phase group that is also contained within every branch local subgroup, we saw that it is possible for the agents to act on separate branches in such a way that both of the constant functions produce the same operational output state.

Finally, when we considered Spekkens' toy hidden variable model, we found that the computation was made possible explicitly by the imposition of an uncertainty principle, which censored the measurement of non-local behavior in the underlying underlying state, in a manner that is similar to the censorship of locality violations arising from hidden variables in Bohmian mechanics.

ACKNOWLEDGEMENTS

The author is very grateful for illuminating discussions and correspondence with Felix Binder, Oscar Dahlsten, Daniela Frauchiger, Nana Liu, Markus Müller, Vlatko Vedral, and Benjamin Yadin. The author is grateful for financial support from the John Templeton Foundation and the Foundational Questions Institute. This research was undertaken whilst the author was funded by the Engineering and Physical Sciences Research Council (UK) at the University of Oxford. Some of concepts in this article have also appeared as part of the author's DPhil thesis [26].

---

[1] **D. Deutsch**. *Quantum Theory, the Church-Turing Principle and the Universal Quantum Computer*. Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences, 400(1818):97–117, July 1985.

[2] **D. Deutsch and R. Jozsa**. *Rapid Solution of Problems by Quantum Computation*. Proc. R. Soc. Lond. A., 439(1907):553–558, 1992.

[3] **D. R. Simon**. *On the power of quantum computation*. In Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on, pages 116–123, November 1994.

[4] **L. K. Grover**. *A fast quantum mechanical algorithm for database search*. In Proceedings of the twenty-eighth annual ACM symposium on Theory of computing, STOC '96, pages 212–219, New York, NY, USA, 1996. ACM.

[5] **P. W. Shor**. *Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer*. SIAM J. Comput., 26(5):1484–1509, October 1997.

[6] **R. Cleve, A. K. Ekert, C. Macchiavello, and M. Mosca**. *Quantum algorithms revisited.* Proc. R. Soc. Lond. A., 454(1969):339–354, January 1998.

[7] **S. Lloyd**. *Quantum search without entanglement.* Physical Review A, 61(1):010301, December 1999.

[8] **A. Zeilinger**. *A foundational principle for quantum mechanics.* Found.Phys., 29(4), 1999.

[9] **L. Hardy**. *Quantum Theory From Five Reasonable Axioms.* January 2001.

[10] **C Fuchs**. *Quantum Foundations in the Light of Quantum Information.* arXiv:quant-ph/0106166v1, 2001.

[11] **Ll. Masanes and M. P. Müller**. *A derivation of quantum theory from physical requirements.* New Journal of Physics, 13(6):63001, 2011.

[12] **G. Chiribella, G. M. DAriano, and P. Perinotti**. *Quantum Theory, Namely the Pure and Reversible Theory of Information.* Entropy, 14(12):1877–1893, October 2012.

[13] **H. Barnum, M. P Müller, and C. Ududec**. *Higher-order interference and single-system postulates characterizing quantum theory.* New Journal of Physics, 16(12):123029, December 2014.

[14] **J. Clauser, M. Horne, A. Shimony, and R. Holt**. *Proposed Experiment to Test Local Hidden-Variable Theories.* Physical Review Letters, 23(15):880–884, October 1969.

[15] **B. S. Cirel'son**. *Quantum generalizations of Bell's inequality.* Letters in Mathematical Physics, 4(2):93–100, March 1980.

[16] **B. S. Tsirelson**. *Some results and problems on quantum Bell-type inequalities.* Hadronic Journal Supplement, 8:329–345, 1993.

[17] **S. Popescu and D. Rohrlich**. *Quantum nonlocality as an axiom.* Found. Phys., 24(3): 379–385, 1994.

[18] **S. Popescu**. *Nonlocality beyond quantum mechanics.* Nature Physics, 10(4):264–270, April 2014.

[19] **W. van Dam**. *Nonlocality & Communication Complexity.* DPhil thesis, University of Oxford, 2000.

[20] **R. Cleve, W. van Dam, M. Nielsen, and A. Tapp**. *Quantum entanglement and the communication complexity of the inner product function.* Theoretical Computer Science, 486: 11–19, May 2013.

[21] **R Sorkin**. *Quantum mechanics as quantum measure theory.* Mod. Phys. Lett. A, 9:3119, 1994.

[22] **C. Ududec, H. Barnum, and J. Emerson**. *Three Slit Experiments and the Structure of Quantum Theory.* Foundations of Physics, 41(3):396–405, February 2010.

[23] **C. Ududec**. *Perspectives on the Formalism of Quantum Theory.* PhD thesis, University of Waterloo, 2012.

[24] **A. J. P. Garner, O. C. O. Dahlsten, Y. Nakata, M. Murao, and V. Vedral**. *A framework for phase and interference in generalized probabilistic theories.* New Journal of Physics, 15(9):093044, September 2013.

[25] **O. C. O. Dahlsten, A. J. P. Garner, and V. Vedral**. *The uncertainty principle enables non-classical dynamics in an interferometer.* Nature communications, 5(4592), August 2014.

[26] **A. J. P. Garner**. *Phase and interference phenomena in generalized probabilistic theories.* DPhil thesis, University of Oxford, Hilary Term 2015.

[27] **R. W. Spekkens**. *Evidence for the epistemic view of quantum states: A toy theory.* Phys. Rev. A, 75(3):32110, March 2007.

[28] **Robert W. Spekkens**. *Quasi-quantization: classical statistical theories with an epistemic restriction.* Quantum Theory: Informational Foundations and Foils, 83-135, December 2015.

[29] **J. Barrett**. *Information processing in generalized probabilistic theories.* Phys. Rev. A, 75(3): 32304, March 2007.

[30] **M. A. Nielsen and I. L. Chuang**. *Quantum Computation and Quantum Information.* Cambridge University Press, Cambridge, October 2000.

[31] **S. J. van Enk**. *A Toy Model for Quantum Mechanics.* Foundations of Physics, 37(10): 1447–1460, August 2007.

[32] **P. Janotta and R. Lal**. *Generalized probabilistic theories without the no-restriction hypothesis.* Physical Review A, 87(5):052131, May 2013.

[33] **Matthew F. Pusey**. *Stabilizer Notation for Spekkens' Toy Theory.* Foundations of Physics, 42(5):688–708, March 2012.

[34] **Y. Aharonov and D. Bohm**. *Significance of Electromagnetic Potentials in the Quantum Theory.* Physical Review, 115(3):485–491, August 1959.

[35] **D. Gross, M. Müller, R. Colbeck, and O. C. O. Dahlsten**. *All Reversible Dynamics in Maximally Nonlocal Theories are Trivial.* Phys. Rev. Lett., 104(8):80402, 2010.

[36] **B. Dakic and C. Brukner**. *The classical limit of a physical theory and the dimensionality of space.* July 2013.

[37] **M. P. Müller and Ll. Masanes**. *Three-dimensionality of space and the quantum bit: an information-theoretic approach.* New Journal of Physics, 15(5):053040, May 2013.

[38] **A. J. P. Garner, M. P. Müller, and O. C. O. Dahlsten**. *The quantum bit from relativity of simultaneity on an interferometer.* December 2014.

[39] **E. M. Alfsen and F. W. Shultz**. *Geometry of State Spaces of Operator Algebras.* Birkhäuser, Boston, 2003.

[40] **L. E. Dickson**. *On Quaternions and Their Generalization and the History of the Eight Square Theorem.* Annals of Mathematics, Second Series, 20(3):155–171, March 1919

[41] **M. A. Graydon**. *Quaternions and Quantum Theory.* Master's thesis, University of Waterloo, 2011.

[42] **J. C. Baez**. *Symplectic, Quaternionic, Fermionic*, 2014. `http://www.math.ucr.edu/home/baez/symplectic.html`.

[43] **D. Bohm**. *A Suggested Interpretation of the Quantum Theory in Terms of "Hidden" Variables. I.* Physical Review, 85(2):166—-179, 1952.

[44] **D. Bohm**. *A Suggested Interpretation of the Quantum Theory in Terms of "Hidden" Variables. II.* Physical Review, 85(2):180–193, January 1952.

[45] **Travis Norsen**. *The pilot-wave perspective on spin.* American Journal of Physics, 82(4): 337–348, April 2014.

Appendix A: Phase and branch locality in quantum theory

First, we show that the phase group of a projective measurement on an N-level quantum system is given by $\bigoplus^N \mathrm{U}(1)$.

Lemma 1. *For an N-level quantum system, with a projective measurement Z, whose outcomes may be expressed in terms of projectors $\left\{Z_j = |j\rangle\langle j|\right\}_{j=1\ldots N}$ (and $\sum_j Z_j = \mathbb{1}$), all reversible phase operations (i.e. satisfying definition 1) may be written in the form*

$$U = \sum_j e^{i\phi_j} |j\rangle\langle j|. \tag{A1}$$

Conversely, all matrices of this form are a phase operation with respect to Z.

Proof. We will use the isomorphism between states s and density matrices ρ, effects e and measurement operators M, such that the inner product in both cases gives the associated probability of a particular outcome ($e \cdot s \leftrightarrow \mathrm{tr}\,(M\rho)$), and the fact that general reversible quantum transformations act as $\rho \to U\rho U^\dagger$ for $U \in \mathrm{U}(N)$. This allows us rewrite the condition for a reversible phase operation with respect to $\{z_i\}$ (written in the GPT framework as $z_j \cdot Ts = z_i \cdot s \quad \forall j, s$) as an equivalent statement in the quantum formalism:

$$\mathrm{tr}\left(Z_j U \rho\, U^\dagger\right) = \mathrm{tr}\,(Z_j \rho) \qquad \forall \rho, j. \tag{A2}$$

This implies that $[U, Z_j] = 0$ for all j, and so we can write U diagonally in the basis of Z,

$$U = \sum_j e^{i\phi_j} |j\rangle\langle j|. \tag{A3}$$

To show the converse, because U of the above form commutes with all elements Z_j, and $U^\dagger U = \mathbb{1}$, we may use the cyclicity of trace to write: $\mathrm{tr}\left(Z_j U \rho\, U^\dagger\right) = \mathrm{tr}\left(U^\dagger Z_j\, U \rho\right) = \mathrm{tr}\left(U^\dagger U Z_j \rho\right) = \mathrm{tr}\,(Z_j \rho)$, which is exactly the phase condition written in eq. (B2). □

From this, we can show the form an operation local to one branch must take agrees with eq. (3).

Lemma 2. *Branch locality implies that the change induced by a reversible quantum operation U_i on a branch i in a set of disjoint branches can be most generally expressed by a complex phase acting only on the ket $|i\rangle$ associated with being in that particular branch:*

$$U_i = e^{i\Phi}\left(e^{i\phi_i}|i\rangle\langle i| + \sum_{j\neq i} |j\rangle\langle j|\right). \tag{A4}$$

Proof. For a branch i, and a state $\rho_{\neg i}$ that has no support on i such that

$$\mathrm{tr}\,(\rho_{\neg i} Z_i) = 0, \tag{A5}$$

then a branch local operation U_i on branch i always satisfies

$$\mathrm{tr}\left(U_i \rho_{\neg i} U_i^\dagger M\right) = \mathrm{tr}\,(\rho_{\neg i} M) \tag{A6}$$

for any observable M. This implies that

$$U_i \rho_{\neg i} U_i^\dagger = \rho_{\neg i}. \tag{A7}$$

The most general state that satisfies eq. (A5) may be written in the Z basis as

$$\rho_{\neg i} = \sum_{\alpha,\beta \neq i} \rho_{\alpha\beta} |\alpha\rangle\langle\beta|. \tag{A8}$$

According to lemma 1, we can express a general quantum phase transformation with respect to Z as $U = \sum_j e^{i\phi_j} |j\rangle\langle j|$, and consider its action on $\rho_{\neg i}$:

$$U_i \rho_{\neg i} U_i^\dagger = \sum_j \sum_{\alpha,\beta \neq i} \sum_k e^{i\phi_j} \rho_{\alpha\beta} e^{-i\phi k} |j\rangle\langle j|\alpha\rangle\langle\beta|k\rangle\langle k| \tag{A9}$$

$$= \sum_{\alpha,\beta \neq i} e^{i\left(\phi_\alpha - \phi_\beta\right)} \rho_{\alpha\beta} |\alpha\rangle\langle\beta|. \tag{A10}$$

Therefore, for the condition in eq. (A7) to be satisfied, we require

$$\sum_{\alpha,\beta \neq i} e^{i\left(\phi_\alpha - \phi_\beta\right)} \rho_{\alpha\beta} |\alpha\rangle\langle\beta| = \sum_{\alpha,\beta \neq i} \rho_{\alpha\beta} |\alpha\rangle\langle\beta|. \tag{A11}$$

For this to hold, it must be true for every element $|\alpha\rangle\langle\beta|$, and this requires $\phi_\alpha = \phi_\beta$ (or for ϕ_α and ϕ_β to differ by an unmeasurable multiple of 2π) for all $\alpha \neq i, \beta \neq i$. A transformation U_i local to branch i thus has the form

$$U_i = e^{i\Phi} \left(e^{i\phi_i} |i\rangle\langle i| + \sum_{j \neq i} |j\rangle\langle j| \right) \tag{A12}$$

where Φ is some global phase that can be ignored.

□

Appendix B: Phase and branch locality in quaternionic quantum theory

By using insights from Graydon [41], we can generalise the results of appendix A into quaternionic quantum theory. A quaternion $h \in \mathbb{H}$ may be written $h = a + ib + jc + kd$ where $a, b, c, d \in \mathcal{R}$ and i, j and k are three types of imaginary number satisfying $ii = jj = kk = ijk = -1$. Conceptually, the jump from complex numbers to quaternions is like the jump from real numbers to complex numbers [40]. Moreover, in the same way that for n-dimension vector-space the orthogonal group $\mathrm{O}(n)$ preserves an inner product between real vectors, and the unitary group $\mathrm{U}(n)$ preserves an inner product between complex vectors, there is a group known as the *symplectic group* $\mathrm{Sp}(n)$ which preserves an inner product for quaternionic vectors [42].

Before we can talk about the quaternionic phase-group, we remark that it is indeed meaningful to talk about quarternionic measurements. Much like in quantum theory, one can find positive semi-definite quaternionic matrices $\{E_i\}$ which act as effects that assign probabilities to on quaternionic states ρ according to $P = \mathrm{tr}\,(E_i \rho)$, and constitute a quaternionic measurement (guaranteed to give one outcome) when $\sum_i E_i = \mathbb{1}$. Moreover one can find sets of N such matrices satisfying $E_i E_j = \delta_{ij} E_i$ which constitute orthonormal projective measurements [41]. Much like the trace on complex matrices, the quarternionic trace is basis-independent, and its arguments may be cyclically rotated [41].

Lemma 3. *For an N-level quaternionic quantum system, with a measurement Z which distinguishes perfectly between N possibilities, the group of diagonal quaternionic matrices*

$$\bigoplus^{N} \mathrm{Sp}(1)\,, \tag{B1}$$

are within the group of reversible phase operations (i.e. satisfying definition 1).

Proof. Take a set of quaternionic projectors $\{Z_i\}$, and consider the most general reversible transformation $S \in \mathrm{Sp}(d)$, which acts on a quaternionic state ρ as $\rho \to S\rho S^\ddagger$ (we shall use $\ddagger$ to represent the symplectic inverse such that $SS^\ddagger = S^\ddagger S = \mathbb{1}$). The phase group condition in definition 1 may be written in terms of quaternionic states, effects and transformations as:

$$\mathrm{tr}\left(Z_j S\rho S^\ddagger\right) = \mathrm{tr}\left(Z_j\rho\right) \qquad \forall \rho, j. \tag{B2}$$

If $[Z_j, S] = 0$ for all j, then this will automatically be true for all ρ. However, we should be careful, as unlike complex matrices, quaternionic matrices do not necessarily commute when they are diagonal (because scalar quaternions do not in general commute), and so we can not automatically conclude that diagonal quaternionic matrices will satisfy this requirement. However, we can use the quaternionic spectral theorem to find a basis where all Z_j are real and diagonal. In this basis the diagonality of S is sufficient to ensure that it can commute with Z_i (because $ah = ha$ for $a \in \mathcal{R}$, $h \in \mathbb{H}$). Hence, we can choose an diagonal matrix of quarternions for S, which in order to satisfy $S^\ddagger S = \mathbb{1}$ must be of the form $\bigoplus \mathrm{Sp}(1)$ (i.e. each diagonal element is a unit quaternion). □

For the discussion in this text, we do not need to claim this is the most general quaternionic phase operation (although it may be so). However, we will show that within this particular group, we can still find non-trivial subgroups of branch local operations.

Lemma 4. *An operation that induces an element of* $\mathrm{Sp}(1)$ *on a branch j, and does not change the other branches except by a global phase of Z_2 is consistent with branch locality as applied to branch j.*

Proof. The proof is similar as with complex quantum theory. For S to be local to branch j with associated measurement effect Z_j, then any state where $\mathrm{tr}\left(Z_j\rho\right) = 0$ must not be changed. For simplicity, we shall use a basis where Z_j is real and diagonal (as is allowed by the spectral theorem).

The density matrix of states with no support in branch j has element $\rho_{jj} = 0$. We can argue that this means all other elements in row j and column j must be 0, using similar logic as holds in the usual quantum case: Pure states in quaternionic quantum theory are rank-1 projectors (Lemma 3.1.1 in [41]), which may be written $|\phi\rangle\langle\phi|$ (for quaternionic bras and kets); and so any pure state with no support in branch j will have $\phi_j = 0$; and as 0 conjugates to itself, this means that the entire row j and column j in the matrix. As mercifully $hh^\ddagger \geq 0$ for $h \in \mathbb{H}$, if a state has no support for branch j, then every pure state it is composed from (which is always possible according to the spectral theorem) must have no support for branch j. Each of these will also only have 0s in row j and column j, and so all of these elements in ρ will be 0.

When S is a diagonal matrix where every diagonal element is 1 except for S_{jj}, branch locality will therefore automatically be respected for density matrices with 0s in row j and column j. Hence, we can associate this choice of $S \in \mathrm{Sp}(1)$ with that branch.

Unlike with quantum theory, the global phase freedom for quaternions can not be the same Sp(1) phase freedom available to each branch: for $h \in \mathrm{Sp}(1)$, application of the global phase $G = \mathrm{diag}\,(h, \ldots h)$ gives $G\rho G^{\ddagger} = h\rho h^{\ddagger}$, which amounts to conjugating every element in ρ with h. Because quaternions do not commute, in general this will have a non-trivial effect on ρ (even if h is restricted to a complex number). Only the real elements of Sp(1), which correspond to $\{1, -1\} = Z_2$, commute and thus are never observable globally. Hence only Z_2 can be taken as global phase operations which have no effect on any observable state (see also Lemma 33 and discussion in Appendix B of [13]). □

This works to our advantage: unlike a complex global phase, a quaternionic global phase would cause us problems with commutativity: if two spatially disjoint agents could both act to span the entire phase-group of a quaternionic two-level system, then relativistic observers might disagree on what the resultant transformation should be, and this would ultimately conflict with the objective reality associated with probabilistic clicks in the detector [38]. Without a quaternionic global phase, branch locality can divide the quaternionic phase group up into subgroups, each of which with elements that commute with all elements in the other subgroups. (The $N = 2$ case is discussed in the Appendix of [38], where the two phase groups end up corresponding to the left and right isoclinic rotations in SO(4), and the only shared element other than identity corresponds to the inversion $1 \to -1$, which is what we require for our interferometric computation.)